

Implementation Guide

Automated Forensics Orchestrator for Amazon EC2



Automated Forensics Orchestrator for Amazon EC2: Implementation Guide

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Solution overview	1
Cost	3
Architecture overview	7
Solution components	10
Forensic triage service	10
Interaction view	10
Implementation view	11
Forensic memory and disk acquisition service	12
Interaction view	12
Implementation view	13
What happens to instances after isolation?	14
Disk forensics acquisition workflow	15
Forensic investigation and reporting service	16
Interaction view	16
Implementation view	17
Forensic investigation and reporting workflow	17
Forensic image and SSM Document builder service	18
Security	20
IAM roles	20
AWS Key Management Service (KMS) Keys	20
Network configuration	20
Data protection	21
Supported deployment Regions	22
Deployment	23
Prerequisites	23
Tools	23
Forensic AMI	23
Compromised instance memory size and investigation instance mount disk volume	24
CDK context configurations	24
Deployment overview	24
Forensic Orchestrator solution deployment in Forensic AWS account	25
Security Hub aggregator account deployment in a new VPC	27
Application account deployment	29
Support for Red Hat Enterprise Linux (RHEL) 8.6 and above	29


Post deployment: Plugin points	30
Memory forensics analysis using LiME and Volatility2	30
Steps to build volatility profile using SSM document	31
Automate the creation of LiME and Volatility 2 profiles	33
Usage of Forensic Solution	33
Step 1. Sign in to the Security Hub AWS Account AWS Management Console and initiate forensic analysis	33
Step 2. Sign in to the Forensic AWS Account AWS Management Console and view Step Functions flow	34
Sample AppSync API to query forensic details	35
Performance considerations	37
AWS services used in this solution	38
Uninstall the solution	42
Using the AWS Command Line Interface (CLI)	42
Using the AWS Management Console	42
Troubleshooting	43
Zero-byte files reported as part of memory and disk investigation	43
ForensicSecHubStack failed to deploy	43
(Optional) Additional configuration Cloud9 environment setup	45
Solution customizations	46
Forensic investigation instance	46
AWS Systems Manager documents	46
Memory acquisition	46
Memory investigation	47
Disk investigation	47
Customization of CDK context configurations	48
Additional resources	53
Source code	54
Revisions	55
Contributors	57
Notices	58
AWS Glossary	59

A self-service solution to capture and examine data from EC2 instances and attached volumes for forensic analysis in the event of a potential security issue being detected

Publication date: *July 2022* ([last update](#): June 2024)

Automated Forensics Orchestrator for Amazon EC2 is a self-service AWS Solution that customers can deploy to quickly set up and configure a forensics orchestration workflow for their Security Operations Center (SOC). It allows their SOC to capture and examine data from EC2 instances and attached volumes as digital forensics evidence for forensic analysis, in the event a potential security breach.


This solution provides a framework to orchestrate and automate key forensics processes from the point at which a threat is first detected. This includes isolation of the affected EC2 instances and data volumes, capture of memory and disk images to secure storage, and initiation of automated actions or tools for investigation and analysis of such artifacts. The solution reports findings and provides process notifications. It allows the SOC to continuously discover and analyze patterns of fraudulent activities across multi-account and multi-region environments. The Automated Forensics Orchestrator for Amazon EC2 solution leverages AWS services and is underpinned by a highly available, resilient, a serverless architecture, security, and operational monitoring features.



Forensic workflow

Digital forensics is a four-step process of acquisition, isolation, investigation and reporting. The Automated Forensics Orchestrator for Amazon EC2 solution provides the capability to act on security events by imaging or acquisition of breached resources for examination and generating a forensic report about the security breach. In the event of a security breach, it allows customers to automatically capture and store targeted data for forensic examination and analysis, and their SOC to discover and analyze patterns of fraudulent activities. The solution supports EC2 instances distributed across multiple accounts and regions.

This solution is intended for deployment in an enterprise by IT infrastructure and security architects, Incident Response team, security administrators, developers, and SecDevOps professionals who have practical experience with the AWS Cloud.

 **Note**

We make no claim as to the suitability of Automated Forensics Orchestrator for Amazon EC2 in the detection or investigation of crime, nor the ability of data or forensics evidence captured by this solution to be used in a court of law. You should independently evaluate the suitability of Automated Forensics Orchestrator for Amazon EC2 for your use case.

Cost

You are responsible for the cost of the AWS services used to run the Automated Forensics Orchestrator for Amazon EC2 solution. As of the recent revision, the monthly cost for running this solution with the default settings in the US East (N. Virginia) AWS Region is approximately **\$235 assuming an average of one forensic instance is 50% utilized for performing forensic analysis** with 512GB of volume attached to the instance. Prices are subject to change. For full details, refer to the pricing page for each AWS service used in this solution.

The total cost to run this solution depends on the following factors:

- The number of forensic incidents reported
- The frequency of forensic orchestration
- The solution assumes a forensic instance runs 12 hours a day

This solution uses the following AWS components, which incur a cost based on your configuration.

Service	Usage estimate	Monthly cost (USD)
AWS Step Functions	Workflow requests (10 per day), State transitions per workflow (20)	\$1
Amazon CloudWatch	Number of Metrics (includes detailed and custom metrics) (20) Number of Custom/Cross-account events (100,000) Number of Dashboards (1) Number of Standard Resolution Alarm Metrics (20) Number of High-Resolution Alarm Metrics (20) Number of Canary runs (5)	\$47

Service	Usage estimate	Monthly cost (USD)
	<p>Number of Lambda functions (10)</p> <p>Number of requests per function (5 per day)</p> <p>Number of Contributor Insights rules for DynamoDB (5) Total number of events for DynamoDB (1 million events per month)</p> <p>Total number of matched log events for CloudWatch (1 million matched log events per month)</p> <p>Number of Contributor Insights rules for CloudWatch (5) Standard Logs: Data Ingested (1 GB)</p> <p>Logs Delivered to S3: Data Ingested (1 GB)</p>	
Amazon DynamoDB	<p>Average item size (all attributes) (20 KB)</p> <p>Data storage size (0.5 GB)</p>	\$26

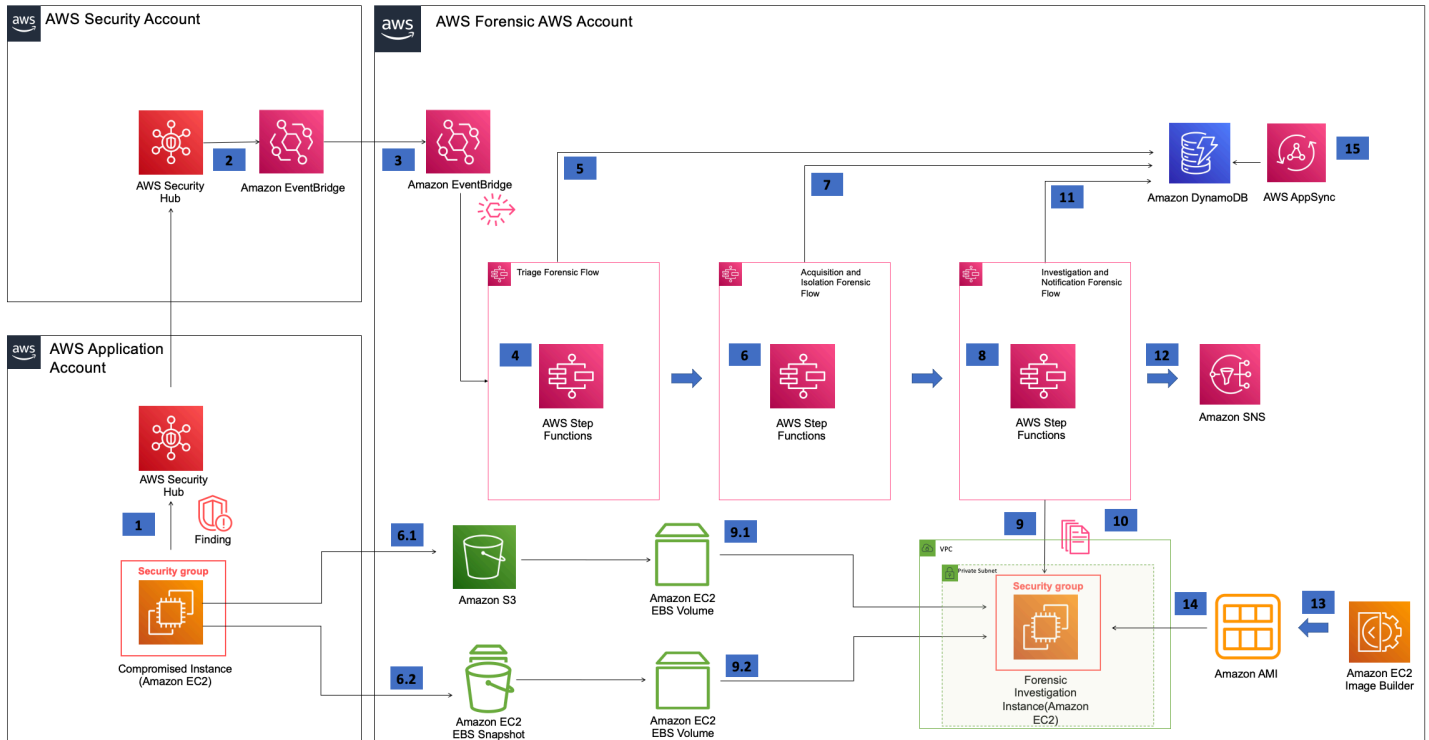
Service	Usage estimate	Monthly cost (USD)
Amazon Simple Notification Service (SNS)	DT Inbound: Not selected (0 TB per month) DT Outbound: Not selected (0 TB per month) Requests (100,000 per month) HTTP/HTTPS Notifications (100,000 per month) EMAIL/ EMAIL-JSON Notifications (100,000 per month) SQS Notifications (100,000 per month) AWS Lambda (1 million per month)	\$2
Amazon EC2*	Operating system (Linux) Quantity (1) Pricing strategy (On-Demand Instances) Storage amount (100 GB) Instance type (M5.2Xlarge) - OnDemand based on the forensic analysis performed – Currently we are leveraging Ubuntu Server 20.04 LTS (HVM) SSD Volume Type	\$142

Service	Usage estimate	Monthly cost (USD)
AWS Lambda	10,000 requests, 60 seconds per lambda function, 512MB of Memory	\$10
AWS KMS Key	1 KMS key, 1,990,000 requests (2,010,000 total requests - 20,000 free tier requests) x \$0.03 / 10,000 requests	\$7
	Total	~\$235 USD/ month

*average usage cost of Amazon EC2

Architecture overview

Deploying this solution with the default parameters builds the following environment in the AWS Cloud.



Automated Forensics Orchestrator for Amazon EC2 architecture diagram

1. In the AWS application account, AWS Config Rule, [Amazon GuardDuty](#), and third-party tools detect malicious activities that are specific to Amazon EC2 resources. For example, an EC2 instance is querying a low reputation domain name that is associated with known abused domains. The findings are sent to AWS Security Hub in the security account via their native or existing integration.
2. By default, all [AWS Security Hub](#) findings are then sent to [Amazon EventBridge](#) to invoke automated downstream workflows.
3. For a specified event, Amazon EventBridge provides an instance ID for the forensics process to target, and initiates the [AWS Step Functions](#) workflow.
4. AWS Step Functions triages the request as follows:
 - a. Gets the instance information
 - b. Determines if isolation is required based on the AWS Security Hub action

- c. Determines if acquisition is required based on tags associated with the instance
 - d. Initiates the acquisition flow based on triaging output
5. Triaging details are stored in [Amazon DynamoDB](#).
6. The following two acquisition flows are initiated in parallel:
- a. *Memory forensics flow* - The AWS Step Function workflow captures the memory data and stores them in [Amazon S3](#). Post memory acquisition, the instance is isolated using security groups. To help ensure the chain of custody, a new security group gets attached to the targeted instance, and removes any access for users, admins, or developers. Note that isolation is initiated based on the selected AWS Security Hub action.
 - b. *Disk forensics flow* - The AWS Step Function workflow takes snapshot of the EBS volume, and shares it with the forensic account.
7. Acquisition details are stored in DynamoDB.
8. Once the disk or memory acquisition process is complete, and the evidence has been captured successfully, a notification is sent to an investigation Step Function state machine to begin the automated investigation of the captured data.
9. Investigation Step Function starts forensic instance from forensic AMI loaded with customer forensic tools:
- a. Loads the memory data from S3 for memory investigation
 - b. Creates an EBS volume from the snapshot and attaches the EBS volume for disk analysis
- 10 AWS Systems Manager documents (SSM documents) are used to run forensic investigation.
- 11 Amazon DynamoDB stores the state of forensic tasks as well as their result when the jobs are complete. Investigation job details are stored in DynamoDB.
- 12 Investigation details are shared with customers using the [Amazon Simple Notification Service](#) (Amazon SNS).
- 13 [EC2 Image Builder](#) builds the Forensic Amazon Machine Images (AMI). Note: You can also use an existing forensic AMI.
- 14 Forensic AMI is leveraged by investigation Step Functions to perform memory and disk investigation.
- 15 The Forensic timeline can be queried using [AWS AppSync](#). For more details, refer to [Sample AppSync API to query forensic details](#).

Note

Using a forensics AMI with the required tooling, and the installed AWS Systems Manager Agent (SSM Agent), the state machine will provision an EC2 instance, attach the previously captured snapshots and mount the memory data captured, making the data ready for investigation. Systems Manager using SSM Run Command runs scripts using the forensic tools installed to perform forensic investigative processes such as timelining against the captured data.

Solution components

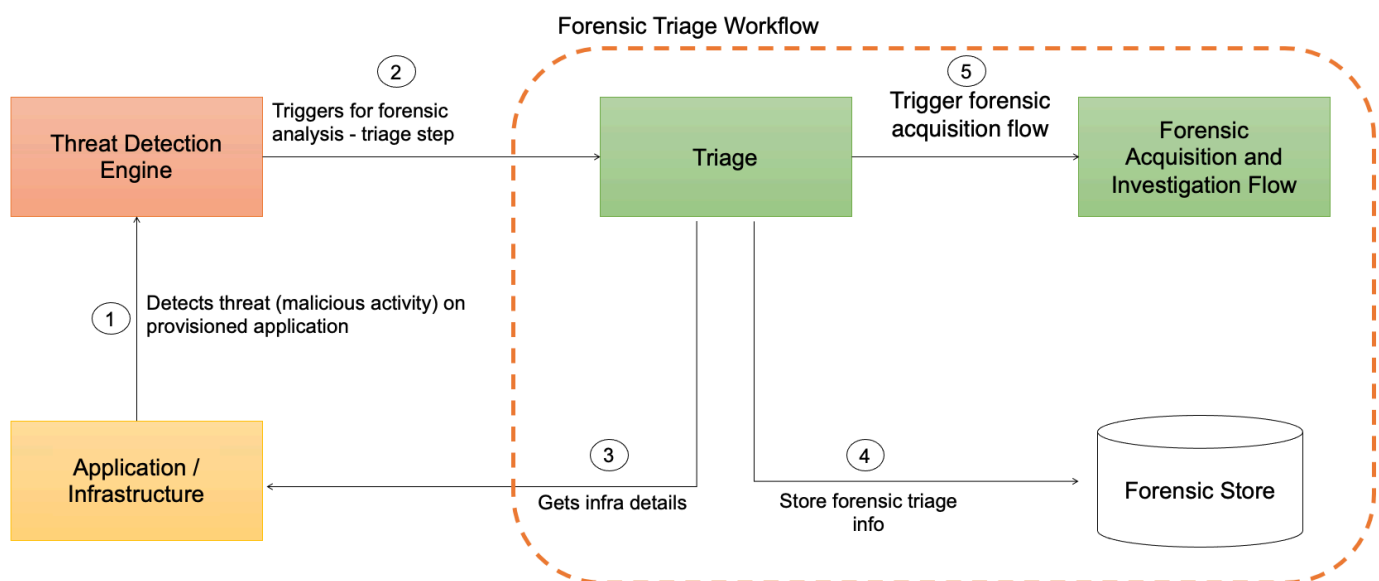
The solution comprises of the following five key components that collaborate to provide EC2 forensic orchestration capability:

- Forensic triage service
- Forensic memory acquisition service
- Forensic disk acquisition service
- Forensic investigation and reporting service
- Forensic image and AWS Systems Manager document builder service

Forensic triage service

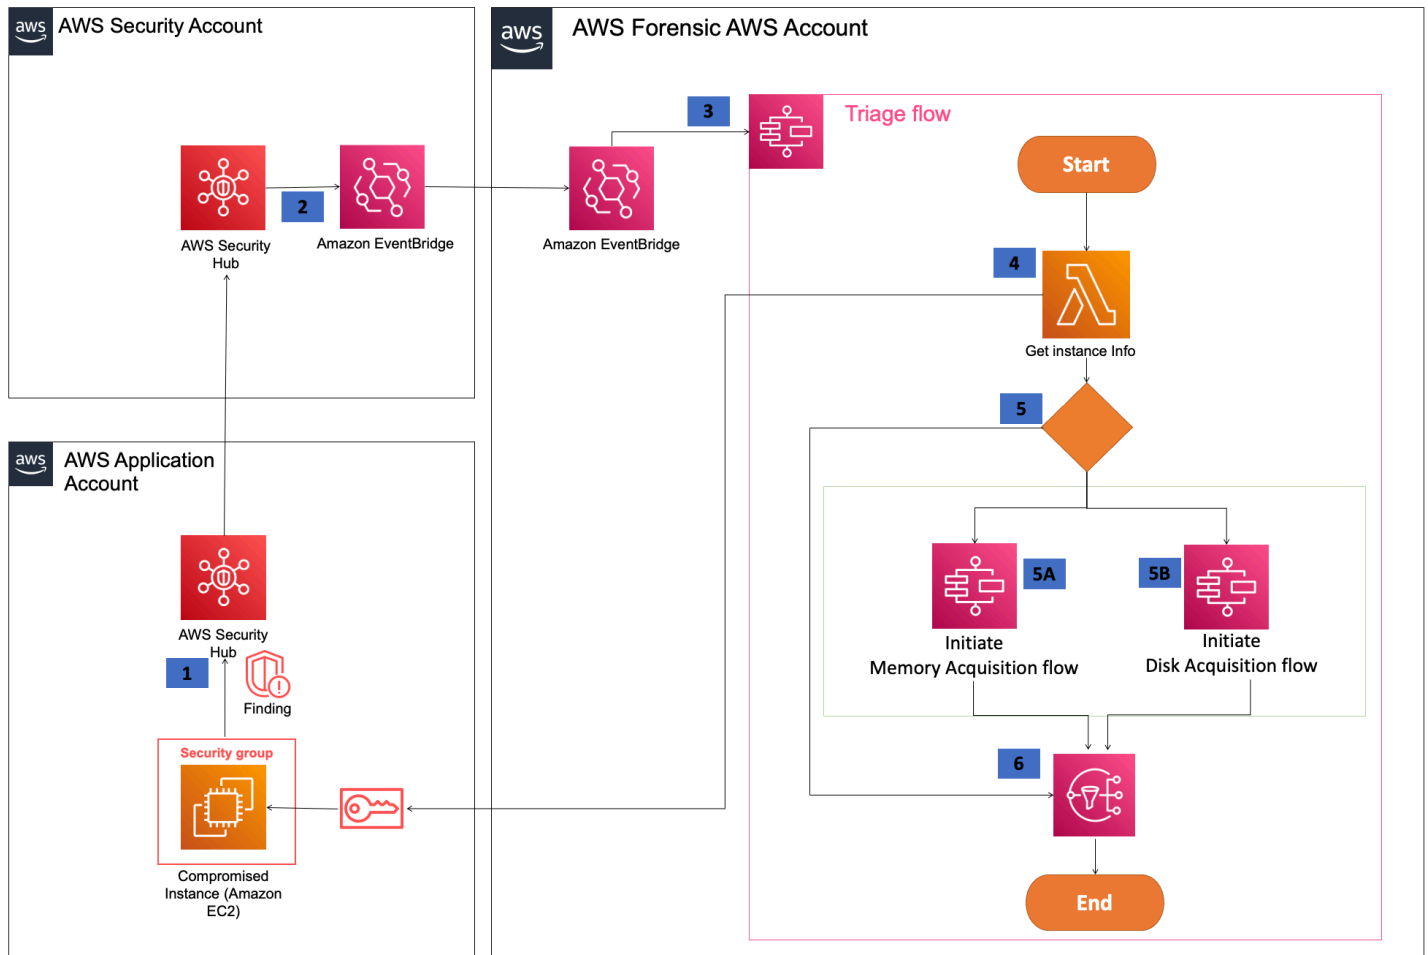
The diagram below represents the logical interaction view of the forensic triage service. A security event (Application security event) is reported by the Threat Detection Engine. The Threat Detection Engine initiates triaging function to determine the severity of threat based on the threat and infrastructure information. The triaging function initiates forensic acquisition and investigation flow for further analysis.

Interaction view



Forensic triage workflow

Implementation view



Forensic triage - implementation view

1. AWS Security Hub operating in AWS application account is reported with details of the compromised instance and the findings get aggregated to AWS Security Hub administrator AWS master Account.
2. The security administrator initiates one of the following forensic actions in Security Hub.
 - a. Forensic triage
 - b. Forensic isolation
3. Amazon EventBridge initiates the *trriage* Step Functions flow.
4. *Get Instance* Lambda function assumes role into compromised application account and retrieves instance information.
5. The *trriage flow* triggers *acquisition flow* in parallel unless the instance tag **IsTriageRequired** is set to `false`.

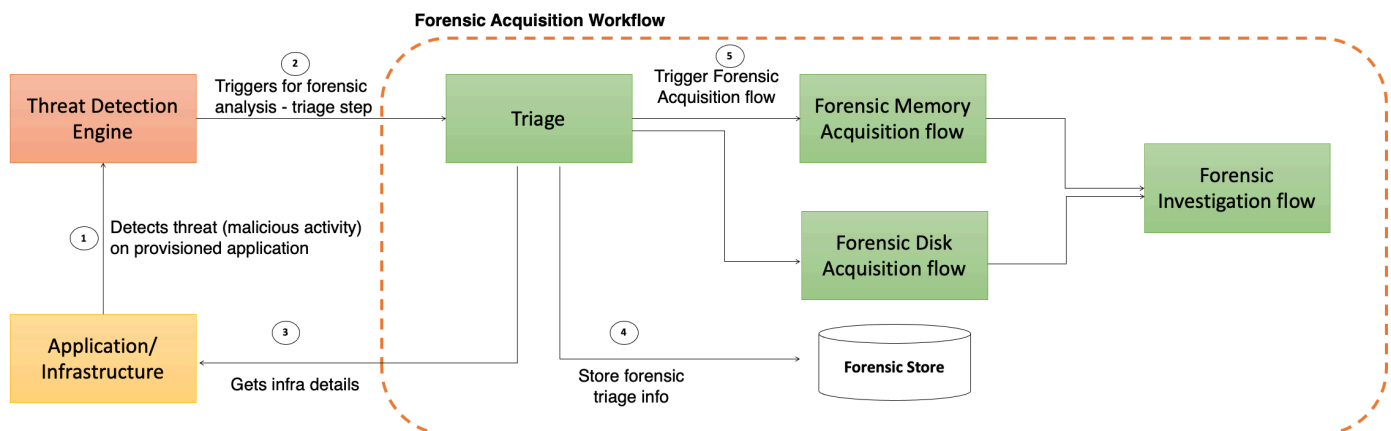
- a. *Forensic memory acquisition flow* initiates the memory acquisition Step Functions.
 - b. *Forensic disk acquisition flow* initiates the disk acquisition Step Functions.
6. Once completed, the acquisition flow triage results are sent to SNS.

Forensic memory and disk acquisition service

The diagram below represents the logical interaction view of the forensic memory and disk acquisition service. The Forensic triaging step function initiates forensic acquisition flow to perform memory and disk acquisition. Following memory and disk acquisition, the investigation function is initiated.

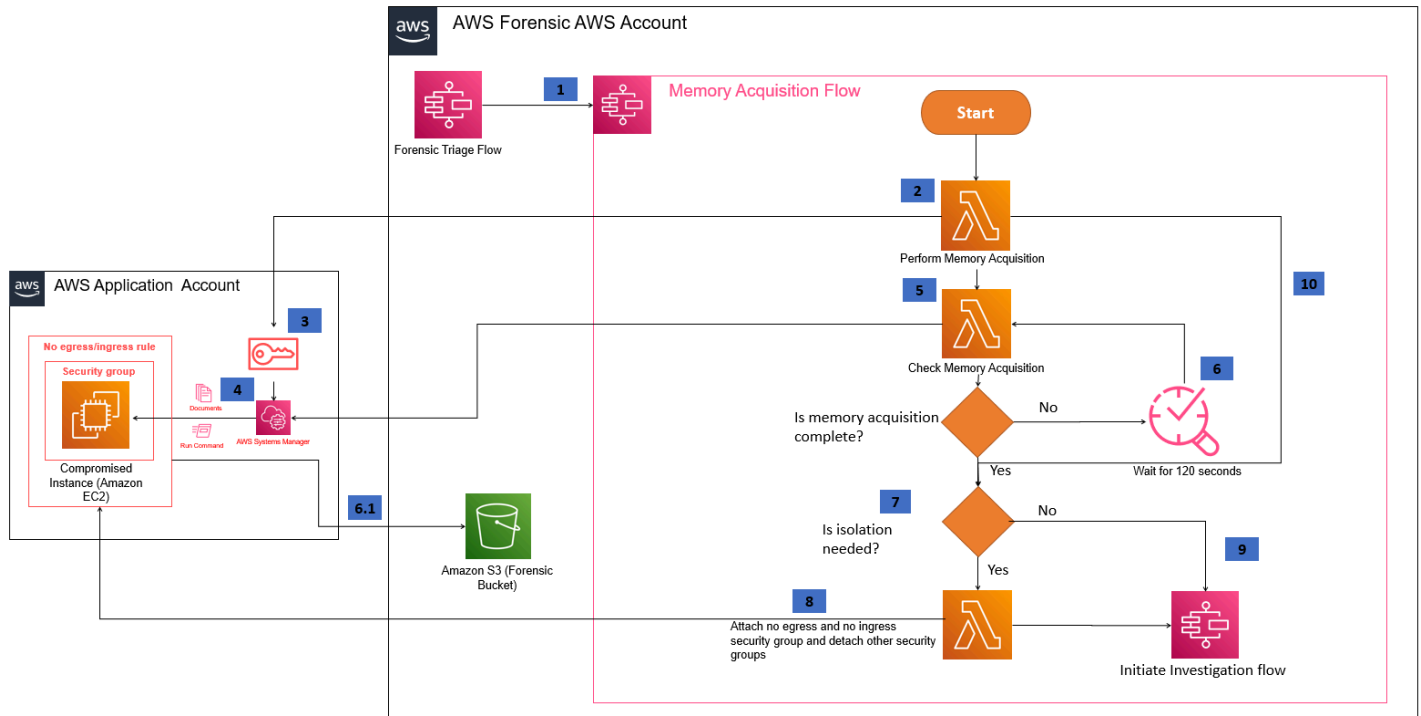
Isolation of EC2 instance is done based on the Security Hub action event types - Forensic triage and Forensic isolation.

Interaction view



Forensic memory disk acquisition - interaction step

Implementation view



Memory forensics acquisition workflow - implementation

1. The *Forensic triage* Step Function initiates the *memory acquisition* flow.
2. The *Memory acquisition* Lambda function in workflow leverages the SSM command to run SSM document in the compromised instance.
3. The *Memory acquisition* Lambda function assumes a role in the application account and passes the SSM document to be run along with credentials to copy the memory dump into an S3 bucket.
4. AWS Systems Manager runs a memory acquisition document via the Run Command.
 - The memory dump is stored in an S3 bucket of the forensic account.
 - The memory dump has associated meta data tags to indicate the underlying OS and kernel the dump is associated with, assisting the *memory analysis flow* further downstream.
5. The *Check memory acquisition* Lambda function checks for SSM Run Command to be completed.
6. If the response from SSM Run Command status is Pending or Delayed or In Progress, it waits for 120 seconds.
7. If the response from SSM Run Command status is Success, it checks if isolation is needed.

8. If `isolation` is set to `true`, then the Lambda function assumes role into the application account and attaches a security group with no egress and ingress security group, and detaches the existing security group. Isolation is set to `true` during the triaging phase based on security event type.
9. This initiates investigation flow with forensic type as `MEMORY`.
10. If any error occurs during the memory acquisition process, the EC2 instance isolation will be performed based on the `isolation` flag.

Note

When the `isolation` flag is set to `true`, isolation is still performed regardless of the memory acquisition result.

What happens to instances after isolation?

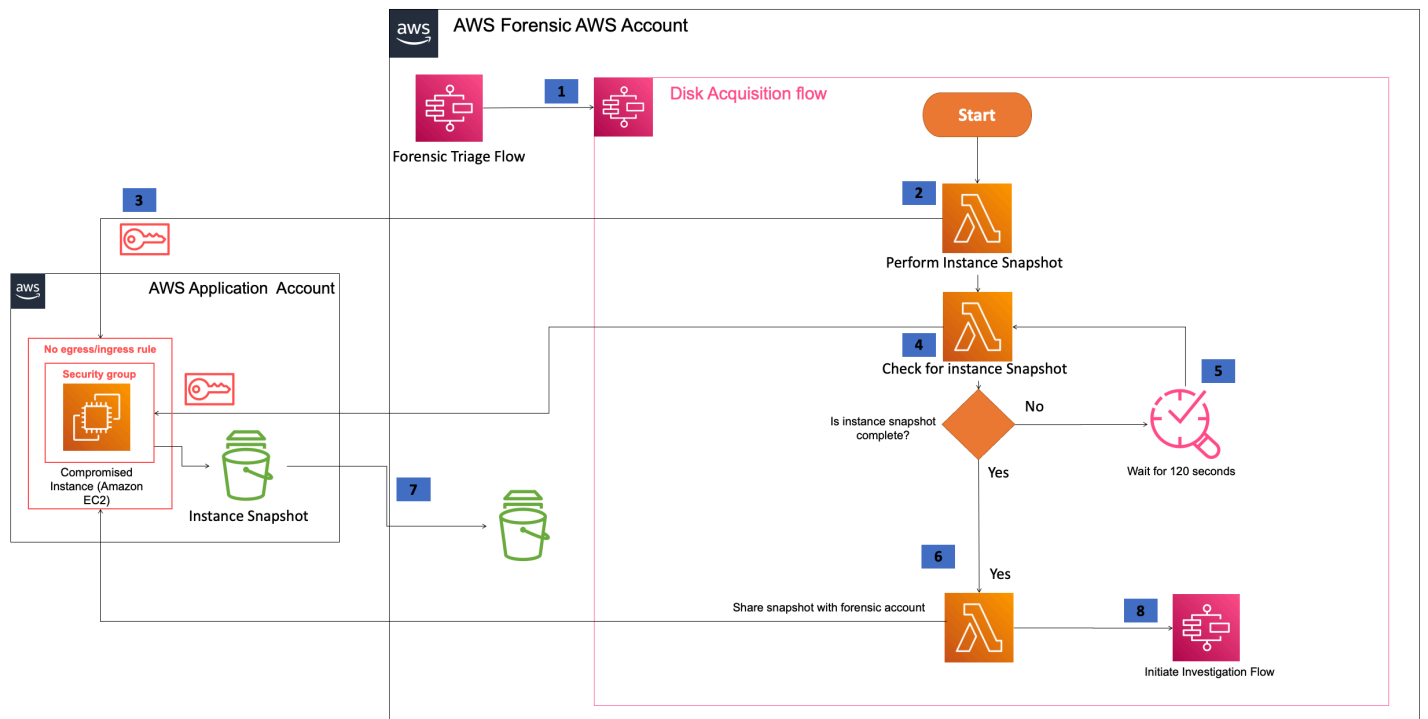
Instances after isolation will have:

- Termination protection for the compromised EC2 instance set to `true`
- Shutdown behavior set to `STOP`
- Any EIP assigned to the compromised instance will be disassociated
- EBS volumes attached to the compromised instance will be preserved
- Instance profile will be updated to a strict profile
- All open credentials session based on the compromised instance role will be invalidated.

Note

On isolation, instances sharing the same role with the compromised instance would be impacted as the credentials will be invalidated. If your application does not have the correct retry mechanism to renew new credentials, it can result in failure of the application. For more information, refer to the [Using temporary credentials with AWS resources](#) topic about creating new credentials. Applications using AWS CLI will not be impacted as the credentials will be refreshed automatically

Disk forensics acquisition workflow



Disk forensics acquisition workflow

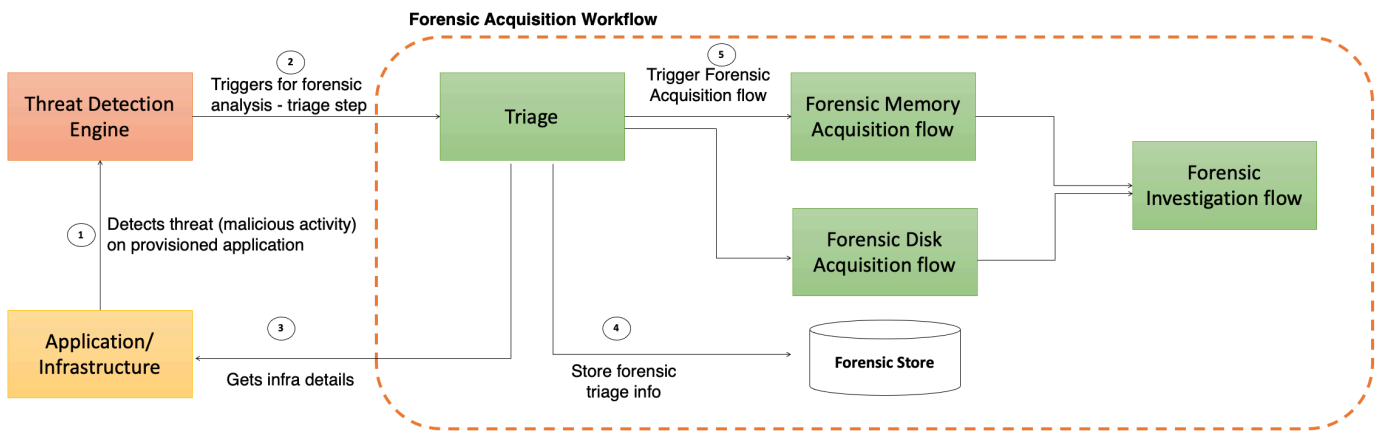
1. The *Forensic triage* Step Function initiates *disk acquisition flow*.
2. The *Perform Instance Snapshot* Lambda function performs an instance snapshot.
3. The *Perform Instance Snapshot* Lambda function assumes a role in the application account and initiates an instance snapshot API call.
4. *Check for Instance Snapshot* Lambda function assumes a role in the application account and checks for snapshot completion.
5. If the response is Pending or In Progress it waits for 120 seconds.
6. The disk acquisition flow [copies](#) the compromised instance snapshot using AWS KMS keys shared with the forensic account.
7. After the copy snapshot operation of the compromised instance, the disk acquisition flow [shares](#) the copied EBS snapshot with the forensic account.
8. To keep the copy of snapshot in the forensic account, copy the shared copy of compromised instance snapshot using Forensic KMS keys. This step allows protection of the shared snapshot with a local copy to perform forensics in case of the AWS account being compromised, or the shared snapshot being deleted by the security team.

9. Post copy flow the step functions initiates the investigation flow with forensic type as DISK.

Forensic investigation and reporting service

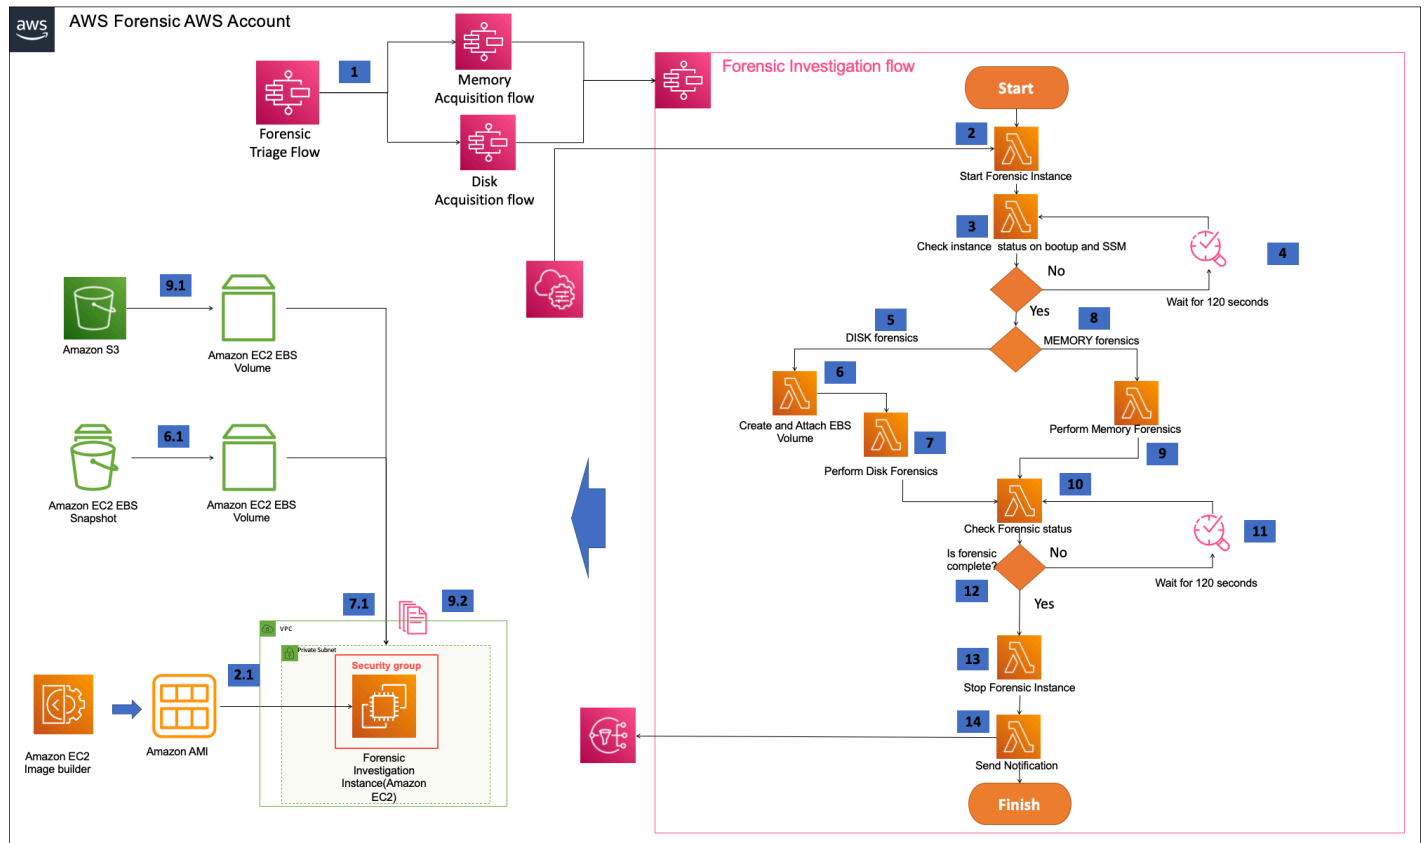
The diagram below represents the logical interaction view of forensic memory and disk investigation service. Once forensic acquisition is completed, forensic investigation flow is initiated, isolation of EC2 instance is done based on the AWS Security Hub event type.

Interaction view



Forensic investigation and reporting service - interaction view

Implementation view



Forensic investigation and reporting service

Forensic investigation and reporting workflow

1. After the acquisition flow, the investigation flow (Step Functions) is initiated.
2. The *Create Instance* Lambda function retrieves the AMI information from AWS Systems Manager Parameter Store and starts an instance in the forensic account.
3. The *Check Instance Lambda* function validates the instance has the necessary tools required for forensic investigation, such as determining the instance is in the running state, AWS Systems Manager is installed and forensic tools are up and running.
4. If the response from SSM Command is Pending or In Progress it waits for 120 seconds and checks again.
5. Disk forensics investigation flow is initiated for **forensictype** variable set to DISK.
6. Disk forensics investigation lambda function creates a volume from the snapshot shared with the forensic account and attaches the volume to the instance started in step 2.

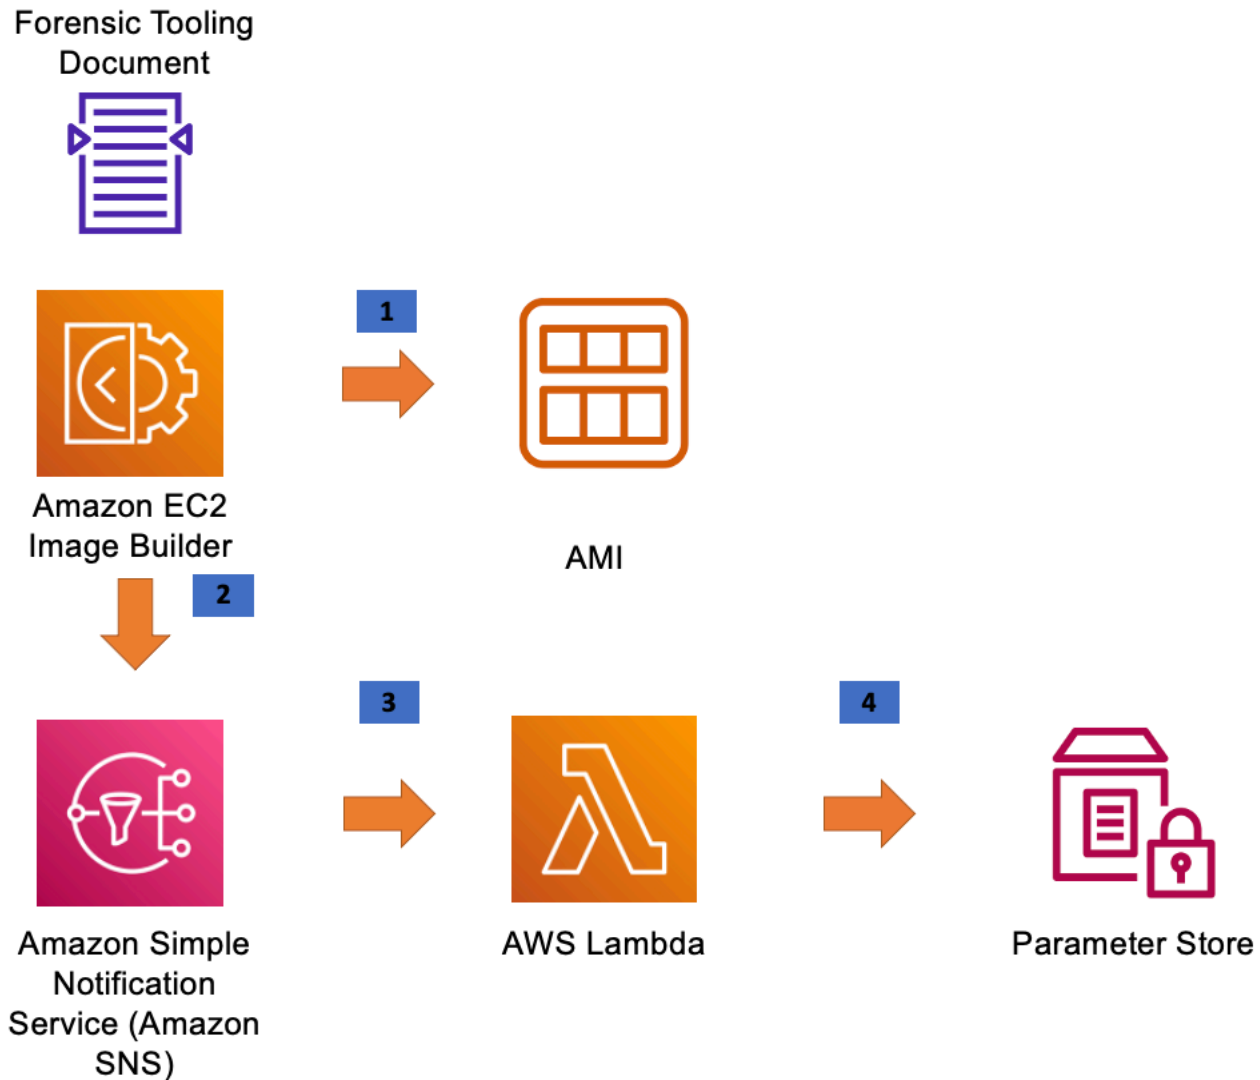
7. The Disk forensics investigation Lambda function leverages the SSM document to perform disk forensics.
8. The Memory forensics investigation flow is initiated for **forensictype** variable set to MEMORY.
9. The Lambda function leverages the SSM document to load memory dump from S3 to the EBS volume for memory analysis.
 - The SSM document containing details of the forensic investigation is initiated to *perform disk or memory forensics*.
 - The *Memory forensics flow* retrieves the appropriate meta data tag associated with the memory dump and loads the matching kernel Volatility profile from a configurable S3 location.
10. The Lambda function checks if AWS Systems Manager Run Command is complete.
11. The Lambda function waits for 120 seconds before checking again if AWS Systems Manager Run Command is complete.
12. Once complete, the *Terminate Forensic Instance* Lambda function is initiated.
13. The forensic instance is terminated.
14. Details about forensic ID, compromised Amazon EC2 instance, Amazon S3 bucket location of the results, and Amazon DynamoDB table details about disk and memory analysis are sent as SNS.

Forensic image and SSM Document builder service

The forensic image builder pipeline creates the forensic AMI with necessary forensic tools needed to perform forensic investigation. The diagram below represents the overall implementation.

Note

Customers can use your own forensic AMI or leverage the [Amazon EC2 Image Builder samples](#) to build a forensic Image.



Forensic image and SSM document builder service

1. Amazon EC2 Image Builder initiates the EC2 Image Builder pipeline to build the EC2 Image based on the forensic tools configured in the document.
2. After successful creation of the AMI, it drops the message as an Amazon SNS topic.
3. The AWS Lambda function listens to the Amazon SNS topic and gets initiated for each message.
4. The AWS Lambda function stores the AMI ID in Parameter Store and is used to launch the forensic instance.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This [shared model](#) reduces your operational burden because AWS operates, manages, and controls the components including the host operating system, the virtualization layer, and the physical security of the facilities in which the services operate. For more information about AWS security, visit the [AWS Cloud Security](#).

IAM roles

AWS Identity and Access Management (IAM) roles allow customers to assign granular access policies and permissions to services and users in the AWS Cloud. This solution creates IAM roles that grant the solution's automated functions access to perform remediation actions within a narrow scope set of permissions specific to each remediation.

AWS Key Management Service (KMS) Keys

The Automated Forensics Orchestrator for Amazon EC2 Framework solution allows you to provide your own AWS KMS keys to encrypt data stored. We recommend referring to [Security best practices for AWS Key Management Service](#) to enhance the protection of your encryption keys.

AWS recommends that customers encrypt sensitive data in transit and at rest. This solution automatically encrypts file data, and metadata at rest with [Amazon S3 Server-Side Encryption \(SSE\)](#) with AES256 algorithm.

Additionally, this solution's Amazon DynamoDB are encrypted at rest using SSE with AWS Key Management Service (AWS KMS).

Network configuration

The Automated Forensics Orchestrator for Amazon EC2 solution is deployed in Amazon VPC, with the Lambda functions in a private subnet. Traffic in and out of the subnet is controlled by security groups. To prevent unauthorized access to the data storage layer, by default, the security group rules only allow inbound traffic from the Lambda function's private subnet.

Data protection

All data committed to Automated Forensics Orchestrator for Amazon EC2 is encrypted at rest, this includes data stored in Amazon S3 and Amazon DynamoDB.

Communications between the solution's different components are over HTTPS to ensure data encryption in transit.

Supported deployment Regions

This solution uses the AWS Step Functions, AWS Lambda, Amazon DynamoDB, AWS EC2 Image Builder, Amazon CloudWatch, Amazon SQS, which are currently available in specific AWS Regions only. Therefore, you must launch this solution in an AWS Region where these AWS services are available. For the most current service availability by Region, refer to the [AWS Regional Services List](#).

Automated Forensics Orchestrator for Amazon EC2 can be deployed in the following AWS Regions in accordance with the regional availability of its constituent services:

Region ID	Region name
us-east-1	US East (N. Virginia)
us-east-2	US East (Ohio)
us-west-1	US West (N. California)
us-west-2	US West (Oregon)
ap-southeast-1	Asia Pacific (Singapore)
ap-southeast-2	Asia Pacific (Sydney)
ap-southeast-3	Asia Pacific (Jakarta)
ap-northeast-1	Asia Pacific (Tokyo)
ap-northeast-2	Asia Pacific (Seoul)
ap-northeast-3	Asia Pacific (Osaka)
ap-south-1	Asia Pacific (Mumbai)
ca-central-1	Canada (Central)
eu-west-1	Europe (Ireland)
eu-west-2	Europe (London)

Deployment

Before you launch the solution, review the architecture, solution components, security, and design considerations discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 30 minutes

Prerequisites

Tools

- The latest version of the [AWS CLI](#) (2.2.37 or newer), installed and configured.
- The latest version of the [AWS CDK](#) V2 (2.2 or newer).
- A [CDK bootstrapped](#) forensic AWS account and Security Hub account.
- [NodeJS](#) version 16.
- Ensure GraphQL – AppSync is activated in the forensic AWS account.
- [AWS SSM agent](#) is installed in EC2 instances (Application Instances).
- AWS Security Hub must be activated as the solution creates custom action in AWS Security Hub.
- Python version 3.8 or above

Forensic AMI

Build a forensic AMI and update the AWS Systems Manager Parameter Store with AMI ID. For more details, refer to [Sample steps to create Forensic AMI using EC2 Image Builder](#). Be sure to replace the image builder component yml with the sample `san-sift.yml` provided in the forensic solution.

Note

This solution requires knowledge of [SAN SIFT](#), [LiME](#) and Volatility tools to customize deployment the solution based on your forensic requirements. For more information, refer to the [the section called “Post deployment: Plugin points”](#) section.

Compromised instance memory size and investigation instance mount disk volume

The investigation instance mount volume must always be greater than the memory of the compromised instance. This ensures that memory loaded into investigation instance does not error out. The solution uses M5.2Xlarge Amazon EC2 instance type by default.

CDK context configurations

CDK config	Description
ec2ForensicImage	Forensic AMI name stored in SSM Parameter Store. SSM Parameter Store will contain AMI ID of forensic investigation instance prebuilt with forensic tools. Solution is tested with Ubuntu AMI

Deployment overview

Use the following steps to deploy this solution on AWS. For detailed instructions, follow the links for each step,

The solution is deployed in the following three AWS accounts:

1. Forensic AWS account – Core solution components to perform forensics orchestration
2. Security Hub AWS account – Configure events and custom actions to trigger forensic orchestration flow
3. Application AWS account – IAM roles needed to establish trust between Forensic AWS account and Application AWS account

Deploying this solution is a three-step process.

1. Forensic Orchestrator solution deployment in the *Forensic AWS Account*.
2. AWS Security Hub configuration to add custom actions to trigger forensics from AWS Security Hub in the *Security Hub AWS account*.

3. *Application AWS Account* deployment to establish trust relationship with the Forensic AWS account.

Note

The Automated Forensics Orchestrator for Amazon EC2 can also be deployed in Security Hub AWS account. Use existing VPC steps to deploy AWS Security Hub configuration in Security Hub AWS account.

Forensic Orchestrator solution deployment in Forensic AWS account

The following steps deploy the Forensics Orchestrator AWS Step Functions, AWS Lambda, and AWS SSM documents into the Forensic AWS account.

1. In your terminal, clone the solution's source code from the [GitHub repository](#).

```
git clone https://github.com/aws-solutions/automated-forensic-orchestrator-for-amazon-ec2.git
```

2. Navigate to the source code folder created in step 1.

```
cd automated-forensic-orchestrator-for-amazon-ec2/source
```

Note

To deploy into existing VPC update `cdk.json` to configure **isExistingVPC** to `true` and add **vpcID** to **vpcConfigDetails** in `cdk.json`.

```
"vpcConfigDetails": {
  "isExistingVPC": true,
  "vpcID": "vpc-1234567890"
  "enableVPCEndpoints": false,
  "enableVpcFlowLog": false
}
```

3. Set AWS credentials to deploy into the AWS account.

```
AWS_ACCESS_KEY_ID=<your_access_key_id>
```

```
export AWS_SECRET_ACCESS_KEY=<your_secret_access_key>
```

```
export AWS_SESSION_TOKEN=<your_session_token>
```

```
export AWS_REGION=<Your Region - us-east-1>
```

4. Install the required NPM libraries.

```
npm ci
```

5. Compile and build AWS Lambda functions.

```
npm run build
```

6. Build the forensics AWS CloudFormation stack to be deployed in the forensic AWS account.

```
cdk synth -c account=<Forensic AWS Account Number> -c region=<Region>  
-c sechubaccount=<Security Hub Aggregator Account Number> -c  
STACK_BUILD_TARGET_ACCT=forensicAccount
```

a. Build the necessary CDK CFN templates for deploying forensic stack. Example:

```
cdk synth -c account=1234567890 -c sechubaccount=0987654321 -c region=us-east-1 -c  
STACK_BUILD_TARGET_ACCT=forensicAccount
```

7. Deploy the forensics stack in the forensic AWS account.

```
cdk deploy --all -c account=<Forensic AWS Account Number> -c region=<Region> --  
require-approval=never -c sechubaccount=<Security Hub Aggregator AWS Account Number>  
-c STACK_BUILD_TARGET_ACCT=forensicAccount
```

Example command that deploys Forensic Solutions stack:

```
cdk deploy --all -c sechubaccount=0987654321 -c  
STACK_BUILD_TARGET_ACCT=forensicAccount -c account=1234567890 -c region=us-east-1 --  
require-approval=never
```

Security Hub aggregator account deployment in a new VPC

As described above, the solution has a dependency on Security Hub to initiate the forensics orchestration. To initiate the forensic Step Functions deployed in the forensic account from AWS Security Hub findings through custom actions present in AWS Security Hub account, deploy the following stack in Security Hub aggregator AWS account.

Note

If you are reusing the existing downloaded code delete the `cdk.out` folder.

1. Clone the solution source code from Solutions GitHub repository.

```
git clone https://github.com/aws-solutions/automated-forensic-orchestrator-for-amazon-ec2.git
```

2. Navigate to the cloned repository created in step 1.
3. Navigate to the source folder.

```
cd automated-forensic-orchestrator-for-amazon-ec2/source
```

Note

To deploy into existing VPC update `cdk.json` to configure `isExistingVPC` to `true` and add `vpcID` to the `vpcConfigDetails` in the `cdk.json` file.

```
"vpcConfigDetails": {
  "isExistingVPC": true,
  "vpcID": "vpc-1234567890"
  "enableVPCEndpoints": false,
  "enableVpcFlowLog": false
}
```

4. Set AWS credentials to deploy into the AWS account.

```
export AWS_ACCESS_KEY_ID=<your_access_key_id>
```

```
export AWS_SECRET_ACCESS_KEY=<your_secret_access_key>
```

```
export AWS_SESSION_TOKEN=<your_session_token>
```

```
export AWS_REGION=<Your Region -us-east-1>
```

5. Install the required NPM libraries.

```
npm ci
```

6. Compile and build AWS Lambda functions.

```
npm run build
```

7. Build the forensics Security Hub AWS CloudFormation stack to be deployed in Security Hub aggregator account.

```
cdk synth -c sechubaccount=<SecHub Account Number> -c  
forensicAccount=<ForensicAccount> -c forensicRegion=us-east-1 -c sechubregion=us-  
east-1 -c STACK_BUILD_TARGET_ACCT=securityHubAccount
```

Example:

```
cdk synth -c sechubaccount=0987654321 -c forensicAccount=1234567890  
-c forensicRegion=us-east-1 -c sechubregion=us-east-1 -c  
STACK_BUILD_TARGET_ACCT=securityHubAccount
```

8. Deploy the forensics Security Hub stack in the Security Hub aggregator account.

```
cdk deploy --all -c sechubaccount=0987654321 -c account=<Security  
Hub AWS AccountNumber> -c region=us-east-1 --require-  
approval=never -c forensicAccount=<Forensic AWS AccountNumber> -c  
STACK_BUILD_TARGET_ACCT=securityHubAccount -c sechubregion=us-east-1
```

Example:


```
cdk deploy --all -c sechubaccount=0987654321 -c account=0987654321 -c
region=us-east-1 --require-approval=never -c forensicAccount=1234567890 -c
STACK_BUILD_TARGET_ACCT=securityHubAccount -c sechubregion=us-east-1
```

Application account deployment

1. Download the [cross-account-role.yml](#) file to your local hard drive.
2. Deploy the `/deployment-prerequisites/cross-account-role.yml` template file as an AWS CloudFormation stack in the application account, and pass the forensic account as input parameter. This will establish a trust relationship between the forensic components deployed in the forensic account and the application account.

```
aws cloudformation deploy --template-file /deployment-prerequisites/
cross-account-role.yml --stack-name app-stack --parameter-overrides
solutionInstalledAccount=<Forensic Solution AWS Account Number>
solutionAccountRegion=us-east-1 kmsKey=<ARN of the application account EBS volume
encryption KMS key>
```

Support for Red Hat Enterprise Linux (RHEL 8.6 and above)

Starting from the current version (1.2.0), the solution supports Red Hat Enterprise Linux (RHEL) 8.6 and above. To use RHEL as the target, you must build a symbol based on the target RHEL.

Note

You will only need to build this once per RHEL version.

To build a symbol:

1. After you deploy the solution, navigate to the [AWS Management Console](#) where you have deployed the solution.
2. Navigate to **Step Functions**, and select the `Forensic-Profile-Function` step function
3. To initiate the build, add these input parameters.

```
{
```

```
"amiId": "ami-0b6c020bf93af9ce1",  
"distribution": "RHEL8"  
}
```

where `ami-0b6c020bf93af9ce1` is the base image Amazon Machine Image (AMI) for RHEL8

Note

You will need a Red Hat subscription before you add this. For more information, refer to the [Linux platforms](#) page.

The `Forensic-Profile-Function` step function will build the symbol automatically. Once the symbol is built, the solution will support RHEL8.

Post deployment: Plugin points

The solution is designed to be an orchestrator and therefore allows the following plugin points where you can replace existing tooling with your preferred tool of choice. The solution leverages [LiME](#) for memory capture and [Volatility2](#) for memory investigation.

1. A LiME module and volatility profile for the EC2 instance must be prebuilt and available for forensic memory investigation and investigation for the EC2 instances OS and kernel version. Refer to Plugin points to build the LiME module and volatility profile.
2. The prebuilt LiME module and volatility profile must be stored in the solution's S3 bucket. The artifacts can be stored in the bucket created by the forensic solution in `cdk.json`. The prefix to the artifacts and the S3 bucket can be configured in `cdk.json` as context variable.

Memory forensics analysis using LiME and Volatility2

As described in [Memory forensics acquisition workflow implementation](#), memory forensics is implemented using Step Functions, which provides the orchestration mechanism and tuns the AWS Systems Manager automation documents. These automation documents can be partially or fully replaced. In the example below, we have bootstrapped our memory forensics implementation to acquire memory using [LiME](#), and to use the [Volatility2](#) profile. LiME is used to extract the volatile memory, which is then analyzed downstream by the Forensics investigation workflow.

When memory is investigated using Volatility 2 (and other tools) it is important for the tool to understand the structure of the memory. In Volatility 2, that is done using a profile which is comprised of two parts: a Debugging With Attribute Record Formats (DWARF) and a map (symbol table used by the kernel) file. The section below demonstrates how the process flow works and ensures that it can be extended by creating a new profile and dropping it into a configurable location (bucket).


Note

As every OS and kernel versions have slight variations, you must create LiME and Volatility artifacts specific to the EC2 instance, and make these artifacts available to the orchestrator during run-time.

In the section below, we walk through how you can build the LiME module and Volatility profile using an SSM document.

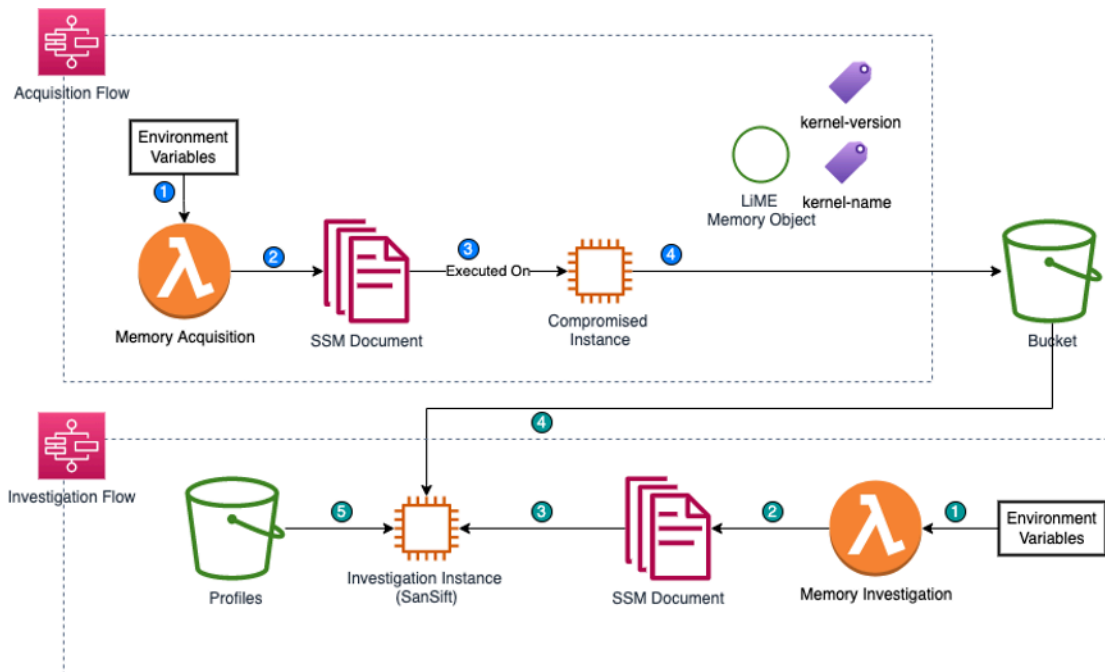
Steps to build LiME module and volatility profile using SSM document

The below diagram explains the overall architecture of building volatility profile.



Steps to build volatility profile using SSM document

The diagram below details the usage of a Volatility profile in the memory investigation flow.



Usage of a volatility profile in the memory investigation flow

1. Launch an [Amazon EC2](#) instance (Amazon Linux 2) to build a LiME module volatility profile. Ensure the SSM is appropriately configured on the EC2 instance. Record the instance ID.
2. Navigate to the AWS Systems Manager documents and select the previously created SSM document example **Documents** tab. Record the name of the SSM document to build the profile.

AWS Systems Manager > Documents > ForensicSolutionStack-ForensicSSMDBuilderStackSSMDocumentamazonlinux2volatilityprofile734DE445-qKoJyi3MMht9

ForensicSolutionStack-ForensicSSMDBuilderStackSSMDocumentamazonlinux2volatilityprofile734DE445-qKoJyi3MMht9

Delete Actions Run command

Description Content Versions Details

Document version: 1 (Default)

Description	Run a script on Linux instances to create profile details	Owner	711045529944
Platform	Linux	Target type	-
Created	Tue, 19 Apr 2022 05:06:15 GMT	Status	Active
Document version	1 (Default)	Latest version	1
Hash type	Sha256	Hash	f219815fa52fb458094a39186d2954c440e6957808d565ba3aabf946127f81d6

SSM document

3. Run AWS SSM document to build LiME module Volatility2 profile for a launched Amazon EC2 instance that matches the OS and kernel version.

Note

Currently the profile and tools are loaded into the S3 bucket for Amazon Linux EC2 instance. For other operating systems, modify the SSM document to create Volatility profiles.

Automate the creation of LiME and Volatility 2 profiles

You can incorporate the module build process for LiME and Volatility (or your preferred forensic tools) into a hardened AMI pipeline prior to allowing AMI use by developers and application teams. These modules are prerequisites for running the Automated Forensics Orchestrator for Amazon EC2 solution to allow the capture and analysis of volatile memory. You also need to incorporate a mechanism to build these modules for the specific kernel versions in the event they do not exist. This can occur if an EC2 instance is updated after being launched or if an EC2 instance was launched from a non-hardened AMI that is not managed by a central team.

For more information, refer to the [How to automatically build forensic kernel modules for Amazon Linux EC2 instances](#) blog, which will walk you through deploying a solution to automatically build modules for specific EC2 instance OS kernel versions based on input parameters of AMI ID and kernel version. You can use the blog solution with the Automated Forensics Orchestrator for Amazon EC2 solution in the event that specific kernel module versions are missing and need to be created.


Usage of Forensic Solution

Step 1. Sign in to the Security Hub AWS Account AWS Management Console and initiate forensic analysis

After the solution CloudFormation stack has been deployed and launched, you can sign in to the web interface.

1. Sign in to the [AWS Security Hub console](#).
2. To display a finding list, do one of the following:
 - a. In the Security Hub navigation pane, choose **Findings**.

- b. In the Security Hub navigation pane, choose **Insights**. Select an insight, and then on the results list, select an insight result.
 - c. In the Security Hub navigation pane, choose **Integrations**. Choose **See findings** for an integration.
3. Select the finding title.
 4. Select the instance findings to trigger forensics.
 5. In **Actions**, you can select:
 - a. **Forensic Isolation** to initiate forensic analysis and perform isolation of instance.
 - b. **Forensic Triage** to initiate forensic analysis.



Forensic analysis - Actions

Step 2. Sign in to the Forensic AWS Account AWS Management Console and view step functions flow

1. Sign in to the [AWS Step Functions console](#).
2. After completion of triaging, an acquisition and investigation flow email will be sent to subscribed SNS topic.



Acquisition and investigation flow - notification

3. Check email for details of the forensic results.

Example Disk Analysis result:

Disk analysis for forensic record b60f5048-2d43-4120-9262-059b3c32bd5f finished successfully. EC2 instance i-0b2cd05ee6445a1c2 in account 930908181379 has been isolated and analyzed. Forensic details are stored in s3 bucket: forensicsolutionstack-forensicbucketawsforensicbu-sgiho6y573vw. For more details on timeline kindly look into Dynamodb table : ForensicTable

Example - Disk Analysis result

Example Memory Analysis result:

Memory analysis for forensic record b60f5048-2d43-4120-9262-059b3c32bd5f finished successfully. EC2 instance i-0b2cd05ee6445a1c2 in account 930908181379 has been isolated and analyzed. Forensic details are stored in s3 bucket : forensicsolutionstack-forensicbucketawsforensicbu-sgiho6y573vw. For more details on timeline kindly look into Dynamodb table : ForensicTable

Example - Memory Analysis result

Sample AppSync API to query forensic details

To query forensic information, [AppSync](#) provides the following queries.

Query	Description
allForensicRecords	<p>Gets all the forensic records. It can be filtered by:</p> <ul style="list-style-type: none"> • awsAccountId • awsRegion • completionTime • creationTime • diskAnalysisStatus • diskAnalysisStatusDescription • id • lastUpdatedTime • memoryAnalysisStatus • memoryAnalysisStatusDescription

Query	Description
	<ul style="list-style-type: none">• resourceId• resourceInfo• resourceType• triageStatus• triageStatusDescription
getForensicRecord	Gets all forensic records based on ForensicID
listForensicRecordsForAccount	Lists forensic records by account.
listForensicRecordsForRegion	Lists forensic records by account and Region.
listForensicRecordsForResource	Lists forensic records by account, Region and ResourceType.
timelineEventsForRecord	Gets timeline of events by ForensicID.

Performance considerations

SSM command timeout

For memory acquisition, memory investigation, and disk investigation SSM documents are leveraged. The timeout is set to 4,000 seconds by default. You can modify the value based on the type of compromised instance. For more details, refer to the following documentation:

- [Handling timeouts in runbooks](#)
- [Understanding command statuses](#)

Compromised instance memory size and investigation instance mount disk volume

The investigation instance mount volume must always be greater than the memory of the compromised instance. This ensures that memory loaded into investigation instance does not error out. The solution uses M5.2Xlarge Amazon EC2 instance type by default.

AWS services used in this solution

Name	Description
AWS Lambda	<p>AWS Lambda is a serverless, event-driven compute service that lets you run code for virtually any type of application or backend service without provisioning or managing servers. You can initiate Lambda from over 200 AWS services and software-as-a-service (SaaS) applications, and only pay for what you use.</p> <p><i>AWS Lambda functions are leveraged to perform forensic actions.</i></p>
Amazon DynamoDB	<p>Amazon DynamoDB is a fully managed, serverless, key-value NoSQL database designed to run high-performance applications at any scale. DynamoDB offers built-in security, continuous backups, automated multi-region replication, in-memory caching, and data export tools.</p> <p><i>Forensic steps are recorded in Amazon DynamoDB for post analysis.</i></p>
Amazon S3	<p>Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. This means customers of all sizes and industries can use it to store and protect any amount of data for a range of use cases, such as data lakes, websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics.</p>

Name	Description
	<p>Amazon S3 provides easy-to-use management features so you can organize your data and configure finely-tuned access controls to meet your specific business, organizational, and compliance requirements. Amazon S3 is designed for 99.999999999% (11 9's) of durability, and stores data for millions of applications for companies all around the world.</p> <p><i>Forensic tools, images and other artifacts are stored in Amazon S3.</i></p>
Amazon CloudWatch	<p>Amazon CloudWatch is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), and IT managers. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events, providing you with a unified view of AWS resources, applications, and services that run on AWS and on-premises servers. You can use CloudWatch to detect anomalous behavior in your environments, set alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to keep your applications running smoothly.</p> <p><i>Forensic orchestration is logged and monitored using Amazon CloudWatch.</i></p>

Name	Description
Amazon EventBridge	<p>Amazon EventBridge is a serverless event bus that makes it easier to build event-driven applications at scale using events generated from your applications, integrated Software-as-a-Service (SaaS) applications, and AWS services. EventBridge delivers a stream of real-time data from event sources such as Zendesk or Shopify to targets like AWS Lambda and other SaaS applications. You can set up routing rules to determine where to send your data to build application architectures that react in real-time to your data sources with event publisher and consumer completely decoupled.</p>
Amazon SNS	<p>Amazon Simple Notification Service (Amazon SNS) is a fully managed messaging service for both application-to-application (A2A) and application-to-person (A2P) communication. The A2A pub/sub functionality provides topics for high-throughput, push-based, many-to-many messaging between distributed systems, microservices, and event-driven serverless applications. Using Amazon SNS topics, your publisher systems can fanout messages to a large number of subscriber systems including Amazon SQS queues, AWS Lambda functions and HTTPS endpoints, for parallel processing, and Amazon Data Firehose. The A2P functionality enables you to send messages to users at scale via SMS, mobile push, and email.</p> <p><i>Forensic notification and reporting is done via Amazon SNS.</i></p>

Name	Description
Amazon EC2 Image Builder	EC2 Image Builder simplifies the building, testing, and deployment of Virtual Machine and container images for use on AWS or on-premises. <i>Amazon EC2 Image Builder can be leveraged to create forensic AMI.</i>
AWS Step Functions	AWS Step Functions is a low-code visual workflow service used to orchestrate AWS services, automate business processes, and build serverless applications. Workflows manage failures, retries, parallelization, service integrations, and observability so developers can focus on higher-value business logic. <i>AWS Step Functions is leveraged to implement Forensic Orchestration flow.</i>

Uninstall the solution

Using the AWS Command Line Interface (CLI)

- Run `cdk destroy --all` from the `sources` folder, or
- Delete the stack from the CloudFormation console in Forensic, Application, and Security Hub AWS Account.

Using the AWS Management Console

1. Sign in to the [AWS CloudFormation console](#).
2. On the **Stacks** page, select this solution's installation stack.
3. Choose **Delete**.

Troubleshooting

Zero-byte files reported as part of memory and disk investigation

Issue: The memory analysis and disk analysis results in S3 could result in zero-byte files.

Name	Type	Last modified	Size	Storage class
timelime.plaso	plaso	April 11, 2022, 23:21:01 (UTC+10:00)	138.0 KB	Standard
vol2-output-linux_bash_sha256.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	93.0 B	Standard
vol2-output-linux_bash.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	0 B	Standard
vol2-output-linux_psaux_sha256.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	94.0 B	Standard
vol2-output-linux_psaux.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	0 B	Standard
vol2-output-linux_pslislist_sha256.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	95.0 B	Standard
vol2-output-linux_pslislist.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	0 B	Standard
vol2-output-linux_psscan_sha256.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	95.0 B	Standard
vol2-output-linux_psscan.txt	txt	April 11, 2022, 23:21:01 (UTC+10:00)	0 B	Standard
webHist.csv	csv	April 11, 2022, 23:21:01 (UTC+10:00)	1.9 KB	Standard

Troubleshooting - zero-byte files

This could be a problem with the Volatility profile. Review the profile associated with the instance as well as the error logs in the Run Command history.

```
Cloning into 'volatility'...
fatal error: An error occurred (404) when calling the HeadObject operation: Key
"volatility2/profiles/5.10.102-99.473.amzn2.x86_64.zip" does not exist
Volatility Foundation Volatility Framework 2.6.1
Volatility Foundation Volatility Framework 2.6.1
ERROR : volatility.debug : Invalid profile selected
```

Troubleshooting - Run Command history

Resolution: This is caused due to an error in SSM document run. Review the SSM error to fix the SSM document. It is necessary to review the Volatility profile and kernel version to ensure it matches the kernel version of running on the compromised EC2 instance.

ForensicSecHubStack failed to deploy

Issue: ForensicSecHubStack failed to deploy. Received response status [FAILED] from custom resource. Message returned: InvalidAccessOperationException: An error occurred (InvalidAccessOperationException)

when calling the CreateActionTarget operation: Account *<Account>* is not subscribed to AWS Security Hub. See details in CloudWatch Log Stream.

Resolution: Activate Security Hub and redeploy ForensicSecHubStack.

(Optional) Additional configuration: Cloud9 environment setup

1. Ensure Python version is 3.8 and above.
2. Set up Python version 3.8 if it is not available (<https://docs.aws.amazon.com/cloud9/latest/user-guide/sample-python.html>).

```
sudo amazon-linux-extras enable python3.8
sudo yum install python3.8
sudo update-alternatives --install /usr/bin/python python
alias python='python3.8'
sudo python -m pip install pip-tools
python -m pip install pip-tools
```

3. (Optional) Update `/source/lambda/requirements.txt` as cloud9 does not support specific versions.

```
arnparse
aws-xray-sdk
boto3
botocore
certifi
charset-normalizer
codeguru-profiler-agent
future
idna
jmespath
jsonpickle
python-dateutil
python2-secrets
requests
s3transfer
six
urllib3
wrapt
```

Solution customizations

Forensic investigation instance

We recommend using M5.2Xlarge instance to perform memory and disk investigation, and initial mount disk volume size as 512 GB to support up to 512 GB of memory. Based on the forensic tools, you can modify the instance type post deployment.

You can update this configuration based on the memory of the compromised instance.

1. [Download the CLI tool to manage a SIFT install.](#)

```
wget https://github.com/teamdfir/sift-cli/releases/download/v1.13.1/sift-cli-linux
wget https://github.com/teamdfir/sift-cli/releases/download/v1.13.1/sift-cli-
linux.sig
wget https://github.com/teamdfir/sift-cli/releases/download/v1.13.1/sift-cli.pub
```

2. Replace the Image Builder document in `san-sift.yml` with S3 copy commands.

Sample S3 command:

```
aws s3 cp s3://mybucket/san-sift
```

Note

The solution uses SANS SIFT to build necessary tools required to perform forensic investigation, and the Amazon EC2 Image Builder to build SANS SIFT image.

AWS Systems Manager documents

The solution uses SSM documents to perform memory acquisition, memory investigation, and disk investigation. You can update the tools leveraged in memory acquisition, memory investigation, and disk investigation with other tools based on forensic requirements.

Memory acquisition

For memory acquisition, the solution provides a [sample SSM document](#), which leverages Linux Memory Extractor (LiME) to perform memory acquisition.

You can update the memory acquisition step with other tools by updating the SSM commands. The sample command clones the GitHub link to set up LiME in the compromised instance. The updated SSM command below downloads the LiME components from an internal S3 bucket to the AWS account.

In SSM document `linux_lime-memory-acquisition.json`, change the following sample command from:

```
git clone https://github.com/504ensicsLabs/LiME
```

to:

```
aws s3 cp s3://mybucket/Lime
```

Memory investigation

You can update the memory investigation step with other tools by updating the SSM commands. The sample command clones the GitHub link to set up Volatility in the compromised instance. The updated command below downloads the Volatility components from an internal S3 bucket to the AWS account. The solution depends on the Volatility profile of the OS and kernel version of the compromised instance. The Volatility profile is a prerequisite to perform memory investigation.

For memory investigation, the solution provides a sample SSM document which leverages Volatility 2 to perform memory investigation.

In SSM document `lime-memory-load-investigation.json`, change the following command from:

```
git clone https://github.com/volatilityfoundation/volatility.git
```

to the sample command:

```
aws s3 cp s3://mybucket/volatility
```

Disk investigation

For disk investigation, the solution provides the sample command leveraging `log2timeline` to perform disk investigation.

You can update the memory investigation step with other tools by updating the SSM commands. The current sample leverages the SANS SIFT image which contains [Plaso](#) in the forensic instance.

Note

The forensic tools code shared in the section are example codes. You are responsible for managing and maintaining the tools used in the solution.

Customization of CDK context configurations

CDK config	Default value	Description
diskSize	512	Determines the disk size of forensic investigation instance. We recommend allocating the disk size more than RAM for the compromised instance.
vol2-profiles-bucket	n/a – leverages the bucket created	Bucket to store Volatility profiles based on kernel version of the compromised instance to perform memory forensics. Ensure forensic investigation IAM role is updated to and has read-only access to S3 bucket.
vol2-profiles-key	/volatility2/profiles/	Prefix in the vol2-profiles-bucket to store the profiles.
customerManagedCMKArns: { forensicsnsEncryptionKey: "" }	n/a	KMS key to encrypt the messages sent to SNS topic.

CDK config	Default value	Description
<code>customerManagedCMKArns:{ forensicBucketEncryptionKey: ""}</code>	n/a	KMS key to encrypt the objects stored in Amazon S3.
<code>customerManagedCMKArns:{ ebsVolumeKey: ""}</code>	n/a	KMS key to encrypt to EBS Volume.
<code>ssmExecutionTimeout</code>	1800	SSM timeout to perform forensic operation. It is set to 1800 seconds by default.
<code>forensicBucketComplianceMode</code>	false	True sets S3 object lock retention mode to compliance else to governance. Refer to S3 Object Lock for more details.
<code>forensicBucketRetentionDays</code>	30	Configures the retention period that protects an object version for a fixed amount of time.
<code>applicationAccounts</code>	*	Contains the list of all accounts the forensic functions can assume role, and perform memory and disk acquisition.
<code>ssm-documents-dir</code>	<code>./ssm-documents</code>	SSM directory in source code to load SSM documents into forensic account as part of the deployment.

CDK config	Default value	Description
ForensicImageName	sansift	Forensic AMI name stored in SSM Parameter Store. SSM Parameter Store will contain AMI ID of forensic investigation instance prebuilt with forensic tools.
ForensicIsolationInstanceProfileName	target_profile_name	Customer profile for isolation . By default it would be the profile provided by this solution's CDK stack.

CDK config	Default value	Description
vpcInfo	<pre>vpcInfo": { "vpcCidr": "10.1.0.0 /16", "maxAZs": 2, "bastionInstance": false, "enableVpcFlowLog": true, "enableVPCEndpoints": true, "subnetConfig": [{ "cidrMask": 24, "name": "externalDMZ", "subnetType": "Public" }, { "cidrMask": 24, "name": "service", "subnetType": "Private" }, { "cidrMask": 24, "name": "database", "subnetType": "Isolated " }, { "cidrMask": 24, "name": "internalDMZ", "subnetType": "Isolated" }] },"</pre>	Contains VPC configurations to create a new VPC.

CDK config	Default value	Description
vpcConfigDetails	<pre>"vpcConfigDetails": { "isExistingVPC": false, "enableVPCEndpoints": false, "enableVpcFlowLog": false }</pre>	<p>Contains details about existing VPC configurations</p> <pre>"isExistingVPC": false, "vpcID": "vpc-1234 567890",</pre>
deployApi	false	<p>Setting it to true deploys GraphQL API.</p> <p>Setting it to false will not deploy GraphQL API.</p>
apiNotifications	false	<p>Setting it to true turns on GraphQL API notification.</p> <p>Setting it to false turns off GraphQL API notification.</p>
apiAllowedIps	[]	Provides list of all IPs allowed to access AppSyncAPI. WAF is configured to restrict the IP address.
apiRateLimit	1000	WAF is configured with <code>ratelimit</code> .

Additional resources

AWS services

- [AWS Identity and Access Management \(IAM\)](#)
- [Amazon Virtual Private Cloud \(Amazon VPC\)](#)
- [AWS Lambda](#)
- [Amazon Simple Storage Service \(Amazon S3\)](#)
- [Amazon DynamoDB](#)
- [Amazon EventBridge](#)
- [Amazon CloudWatch](#)
- [AWS X-Ray](#)
- [Amazon Simple Queue Service \(Amazon SQS\)](#)
- [AWS Step Functions](#)
- [AWS Security Hub](#)
- [Amazon GuardDuty](#)

Source code

Visit our [GitHub repository](#) to download the source files for this solution and to share your customizations with others. The Amazon EC2 Forensic Orchestrator templates are generated using the [AWS Cloud Development Kit \(AWS CDK\)](#). Refer to the [README.md](#) file for additional information.

Revisions

Date	Change
July 2022	Initial release
December 2022	<p>Release v1.1.0</p> <ul style="list-style-type: none">• Added solution customization option• Updated memory isolation flow section with isolation behavior changes.• Added LiME module information to Deployment section (Prerequisites, Post deployment: Plugin points)• Updated Memory forensics analysis using LiME and Volatility 2 and Steps to build LiME module and volatility profile using SSM document sections
May 2023	Added support for Red Hat Linux 8.6 (and above).
July 2023	Mitigated impact caused by new default settings for S3 Object Ownership (ACLs disabled) for all new S3 buckets. For more information, refer to the CHANGELOG.md file in the GitHub repository.
April 2024	<p>Release v1.2.2</p> <p>Removed the metric collector component; and updated dependencies. For more information, refer to the CHANGELOG.md file in the GitHub repository.</p>
June 2024	Release v1.2.3

Date	Change
	Updated package versions to resolve security vulnerabilities. For more information, refer to the CHANGELOG.md file in the GitHub repository.

Contributors

- Deenadayaalan Thirugnanasambandam
- Yang Yang
- Hafiz Saadullah
- Tim O'Hare
- Barry Conway
- Ruskin Dantra
- Jason Martin
- Jonathon Poling
- Jonathan Nguyen
- Iqbal Umair
- Swapnil Ogale
- Daniil Millwood
- Verinder Singh
- Abe Wubshet

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Automated Forensics Orchestrator for Amazon EC2 is licensed under the terms of the of the Apache License Version 2.0 available at [The Apache Software Foundation](https://www.apache.org/licenses/LICENSE-2.0).

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.